For Your Own Epistemic Good

Jordan Myers

5/7/26 Draft

"It's easier to fool people than to convince them they've been fooled"

—Mark Twain

My partner recently asked if I could help intervene on her mother's obsession with conspiracy theories. My mother-in-law has always had an affinity for alternative explanations, but over the past few years, it's reached a new level.[1] She's gone all-in on theories like the 'PizzaGate' scandal, 5G causing cancer, and the idea that no one has ever actually been to the moon. If there's a conspiratorial explanation for any event, she'll buy it.

I spent some time thinking about how to even begin breaking her out of her echo chamber.[2] Clearly, I couldn't just start talking with her naturally, like I would with a close friend or colleague. I couldn't openly report my thoughts on her poor inquiry processes or dispositions to believe anything with a conspiratorial spin; my honest thoughts would immediately brand me an 'insider' or a 'sheep.' Those trapped within echo chambers are highly resistant to any interventions by outside voices (Nguyen 2020, 145-150 & Jamieson and Cappella 2008). I'd instead have to adopt a more strategic view, figuring out what would be most effective and then tactically employing that strategy.

I might try to 'out conspiracy' her, basically attempting to show the unfalsifiable status of her theories. So, when she brings up 5G cell towers causing cancer, rather than simply give my actual thoughts on why that's not likely to be the case, I'd instead pursue a tactic that doesn't align with my actual thoughts. I might try to flood her with even more elaborate theories in an effort to push her towards declaring epistemic bankruptcy (as something like a first step towards recognizing that she is in an epistemic community that has a misfiring disposition to read *everything* as a conspiracy). So maybe I'd suggest that not only do 5G towers cause cancer, but that this fact is only part of a larger, global conspiracy; 5G causes cancer *and* it's blasting our minds with a high frequency signal designed to lower our IQ and make us easier to control. Or maybe I'd suggest that cancer itself isn't real; it's just a liberal hoax of some sort. I could spin up a few more theories like this in an effort to make her realize that her conspiratorial way of thinking is an unconstrained and unfalsifiable methodology. The idea would be to push such implausible theories that even she couldn't accept them.

Regardless of what specific tactic I might pursue, I would attempt, not to argue with her, but simply to persuade her. In essence, I would take Twain's advice. I'd try to fool her rather than convince her she had been fooled. Only here, my aim would be to fool her into believing the right things for the right reasons.[3]

---

[1] At the time of writing, she is not yet my mother-in-law, but I'm being optimistic about both the publication of this paper and the success of my relationship with her daughter.

[2] Nguyen has recently provided an excellent discussion of the differences between an epistemic bubble and an echo chamber: "An epistemic bubble is a social epistemic structure in which other relevant voices have been left out, perhaps accidentally. An echo chamber is a social epistemic structure from which other relevant voices have been actively excluded and discredited" (2020, 141).

[3] Interestingly, Levy has recently argued that conspiratorial thinking is not subjectively irrational and may even be objectively rational, since it results from epistemic dispositions that are themselves rationally defensible (2019). While I disagree with Levy's thesis, I do not have the space to address it here. I'm working under the assumption that my mother-in-law's pornographic conspiracy theorizing is both subjectively and objectively irrational. As will become clear,

There's an interesting question about what exactly I'd be doing when I engaged my mother-in-law in this manner. I wouldn't *simply* be lying. But it also doesn't seem like I'd be deceiving her in any straightforward manner either. I'm trying to fool her in a sense, yes; but I'm not actually saying anything that's dishonest—I actually don't believe any of these theories are true and I don't think that her style of inquiring is rational. Even if I am being dishonest in some sense, this isn't the full explanation of what I'm doing. It's not *just* that I'm discounting her testimony (although I am also doing that), and it it's not *just* that I don't take her to be my epistemic peer.[4] I'm doing something more than these explanations; I'm not taking her seriously as an epistemic agent. I'm being paternalistic in some way towards her, treating her like a child. But paternalism tends to involve interfering with what another person *does*, not what they believe. So is paternalism the right lens through which to view my interference? Partially.

Paternalism is essentially the practice of treating someone like a child. When that person *is* a child, there's nothing wrong with being paternalistic. It's right to restrict what a ten-year-old can do for their own good. Tension arises when we treat other adults in this way. Paternalism has been broadly understood as the act or practice of interfering with an agent's autonomy, without their consent, for that person's own good (Mill 1869, Dworkin 2010, Grill & Hanna 2018).[5] We engage in paternalism with great frequency; the state, for instance, passes seatbelt laws and mandatory social security contributions. Parents engage in overt paternalism towards their children—sometimes far too long into adulthood.

But it's a further question whether and with what frequency we engage in what might be called *epistemic paternalism*: roughly, the practice of "interfering with agents, without their consent, for their own epistemic good—e.g. to promote their true beliefs, knowledge, etc." (Jackson 2021, 132). Do we paternalistically manage others' beliefs in the same way we manage their behavior? How common is this practice? What does it involve and what could justify treating others in this way?

In this paper, I develop a revised account of epistemic paternalism and provide a rule of thumb for differentiating *epistemic* from *non-epistemic* paternalism. This is the goal of Section 1: to explore the concept and offer several revisions to Elizabeth Jackson's recent account. I'll then argue in Section 2 that, contrary to recent descriptive skepticism, epistemic paternalism (EP) is a widespread interpersonal practice. Having argued that EP constitutively involves a lack of respect for another's epistemic capabilities, I'll say more about what it means to respect (or fail to respect) another's epistemic capabilities in Section 3. After this conceptual work, I'll develop a framework for answering normative questions about when and why epistemic paternalism might be justified. This aim comprises Section 4. Finally, in Section 5, I pursue my final aim: to illuminate the paradigmatic mental state involved in epistemic paternalism.

---

however, I don't need this assumption for my central thesis to hold, since epistemic paternalism can be undertaken for only what the paternalistic party *takes* to be the other's good. So I could be wrong in this case and still be right about my central claims in this article.

[4] Whether I take my mother-in-law seriously as an epistemic agent is not precisely the same thing as taking her to be my epistemic peer: "agents who are similarly or equally well-qualified to opine upon matters in a given domain" (Simpson 2013, 563). These two notions come apart in both directions; I can take someone seriously and show them respect for their epistemic capabilities while explicitly thinking they are *not* my epistemic peer, and I can believe someone is my epistemic peer while treating them as though they are not.

[5] There's a distinction between what is called 'weak' vs 'strong' paternalism (Dworkin 2020). Weak paternalism involves interfering with others only when their local action undercuts their own global aims, while strong paternalism takes as legitimate interfering with others even when this cuts across their own global aims or ends. I'll be discussing epistemic paternalism as encompassing both the weak and strong sense.

Section 1: What is Epistemic Paternalism?

As the name implies, epistemic paternalism is a subset of paternalism more broadly. General paternalism is the kind we know as citizens who are subject to laws or as either side of a parent-child relationship. It involves interfering with the realm what other agents *do*. Epistemic paternalism involves interfering with the realm of what other agents *believe*. More accurately, epistemic paternalism involves interference that targets another's epistemic life: their doxastic states or inquiry more broadly.

Following Pritchard (2013) and Bullock (2018), Jackson has recently developed an account of epistemic paternalism. In this section, I will examine her account and offer several critical revisions that I believe are necessary to fully understand the phenomenon. Jackson begins her own investigation by examining a widespread account found in Goldman (1991) and Lougheed (2021). This account is summarized by Jackson:

> "**Epistemic paternalism** = (i) withholding evidence from someone, (ii) without their consent, (iii) to make it more likely that they believe truths (or avoid errors)" (2021, 133).

Jackson rightly highlights several issues with this account. As it stands, condition (iii) is too restrictive: it assumes that "veritism, the view that believing truth (and avoiding error) is the only final epistemic good" (2021, 133). Instead, we should leave open what counts as the aim of epistemic paternalism. Secondly, EP can be pursued without withholding evidence, as condition (i) currently suggests. It seems clear, for instance, that I could paternalistically manage your beliefs by either lying or reporting true facts in a manipulative manner.[6] Condition (i) should therefore refer to the interference of someone's epistemic activity more broadly (2021, 133-135).[7] Finally, the paternalistic party should interfere *intentionally* and *significantly*. This modification aims to prevent rendering cases of ordinary collective reasoning as paternalistic (2021, 135-137). This revision also makes epistemic paternalism an intentional or non-accidental activity; it requires a mental component on behalf of the paternalistic party. One must act paternalistically and *intend* to act this way. With these revisions in hand, Jackson presents her own account:

> "**Epistemic paternalism** = (i) intentionally and significantly interfering with someone's inquiry, (ii) without their consent, (iii) for their own epistemic good" (2021, 137).

As it stands, this account is an improvement. However, I suggest there are three lingering problems.[8] The first issue is minor: this account states that epistemic paternalism is performed for the other's own epistemic good. But surely, one could interfere for what they mistakenly believe to be the right epistemic ends, such that they act paternalistically only towards what they *take* to be the other's own good.[9] The paternalistic party could be right or wrong about what is in fact good for the other person. Thus, we should revise condition (iii):

> **Epistemic paternalism** = (i) intentionally and significantly interfering with someone's inquiry, (ii) without their consent, (iii) for *what the paternalistic party takes to be* their own epistemic good.

Still, this revision is not enough. Two related problems remain:

---

[6] See Bullock (2016 & 2018) for a discussion of paternalism by providing evidence in medical cases and Bansback et al. (2014) for a discussion of how ordering effects influence patients' decisions about medical treatment options.

[7] These changes result in an account similar to that found in Ahlstrom-Vij (2013) and Jackson (2020).

[8] As Jackson herself admits, "this definition is still not sufficiently precise" (2021, 137).

[9] So we still wouldn't call manipulating someone's beliefs for your own gain paternalistic; that's just selfishly manipulating them. The paternalistic party must *believe* that they are acting for the other's own good, even if they end up being wrong about that.

The Over-Generation Problem: This account remains too broad. As it currently stands, even this revised account will render instances of simply sharing information as cases of epistemic paternalism.[10] We often share information, evidence, or reasons with others (intentionally and significantly altering their inquiry process), without their consent, and for their own good (we want people to avoid errors after all). But simply sharing evidence or reasons (even without first checking that it's okay) does not necessarily amount to epistemic paternalism.

Missing Wrong-Making Feature: It's not yet clear what feature generates the moral wrong providing a (pro tanto) reason *against* being epistemically paternalistic. But of course, a staple assumption of the literature is that EP is indeed a (pro tanto) moral wrong.[11] Even among its defenders, it's widely accepted that being epistemically paternalistic is a wrong that, in some cases, can be outweighed by other considerations.[12] But even on this revised account, the wrong-making feature of epistemic paternalism is not yet clear.

One might argue that condition (ii) supplies the source of the necessary moral wrong. Because the intervention is performed without the other's consent, perhaps this provides our missing wrong-making feature. But consent is generally understood as the waiving of a claim-right against others.[13] And it's not at all obvious that we have claim-rights against others simply sharing information, evidence, or reasons with us without our consent. Thus, this avenue for generating the pro tanto moral wrong of epistemic paternalism falls flat.

Likewise, the wrong-making feature cannot simply be found in the *interference*—condition (i). This certainly seems to be part of the pro tanto wrong of general paternalism; people have a wide swath of non-interference rights which paternalism infringes. But again, this explanation only seems plausible for non-epistemic paternalism. Interfering with what people are free to *do* is something that clearly requires justification, but it's not at all clear that simply interfering with others' inquiry is inherently wrong in the same way. I'll illustrate this point in the upcoming pair of cases.[14]

Perhaps the wrong-making feature is engaging in deception or dishonesty. The problem, however, is that neither are necessary to engage in epistemic paternalism. First, one can strategically report the truth in a way that manages another's beliefs without engaging in deception. Secondly, one could also simply withhold true information in an effort to manage another's beliefs or inquiry; a teacher might intentionally *not* provide an answer to her student, forcing him to arrive at the result himself. So the essential wrong-making feature must be located elsewhere.

In sum, we have two remaining issues: our account (1) remains too broad and (2) lacks a wrong-making feature. I believe that one modification can solve both these problems. Consider the following pair of cases:

> ***Future Engineers*:** Two friends, Tom and Jerry, are in the university library working on their individual engineering projects. Each student is tasked with designing a bridge that will be modeled and stress tested. Tom, being better at statics, notices that Jerry has incorrectly

[10] Jackson herself discusses this issue but fails to sufficiently guard against this possibility (see 2021, 136-137). Thanks to Eleanor Gordon Smith for pointing this out.

[11] See Jackson (2021), Levy (2021). Anderson (1999) takes as a foundational assumption that paternalism is a morally bad feature of any political system of governance. Larmore (1990) and Nagel (1987) both take as a basic premise that the coercion involved in paternalism is one way of treating others as mere means and not as ends, and thus must be justified (or normatively overcome).

[12] Critics of nudging, for instance, argue that nudges are wrong *because* they manipulate others' beliefs, which is a common way to be epistemically paternalistic; see (Bovens 2008; Wilkinson 2013; Saghai 2013). Levy (2021) adopts this assumption even when defending the permissibility of nudging. He argues that nudges are permissible because they are *not* threats to our autonomy because they do *not* paternalistically manipulate us.

[13] See Owens (2011), Alexander (2014), Manson (2016), Tadros (2016), Dougherty (2021), and Liberto (2021).

[14] Thanks to Yoorim In for pushing me to emphasize this point.

calculated the moment of inertia for a crossbeam. He then points out this blunder to Jerry (without first asking if he can do so).

This is a normal interaction between peers. But it seems like even our revised account could plausibly render this an instance of epistemic paternalism: Tom intentionally and significantly interfered with Jerry's inquiry process, without Jerry's consent, and did so for what Tom took to be Jerry's own epistemic good (he was wrong after all!). But epistemic interventions like this don't seem like the kind of thing we want our account to capture. Tom isn't being paternalistic here—he's just pointing out a mistake. We want epistemic paternalism to require something beyond merely presenting information or reasons without the recipient's consent, even if this is done intentionally and as a significant interference into the other's inquiry process for their own good. Consider a different case:

> ***Postmillennial Grandfather.*** John wants to convince his grandfather of what he takes to be an objectively true normative claim: that climate change is a problem worth worrying about and solving. In the past, his grandfather has espoused apathy about the issue because, in his own words, '*Jesus will be returning soon anyway, so it doesn't matter if the Earth is a few degrees warmer a hundred years from now.*' John recognizes that his grandfather's stance on climate change is tied to his religious beliefs, and the likelihood that John will be able to undermine his grandfather's religious commitments is exceedingly slim. So, being a crafty epistemic agent, John pursues a different strategy. He knows that his grandfather subscribes to postmillennialism—the doctrine that Christ will return only after a thousand-year golden age of peaceful Christian rule on Earth. John himself thinks all this is rubbish. But he strategically pushes the following argument: If climate change is left unchecked, it will cause massive disruption and chaos on Earth. This will prevent the possibility of the thousand-year golden age of peace, which will prevent Jesus from returning. That would be bad. Therefore, we should work to mitigate climate change. The argument, to John's surprise, works. His grandfather now believes that climate change should be mitigated.

What's the difference between Tom's stance towards Jerry and John's stance towards his grandfather? The first does not at seem paternalistic, while the latter definitively does. Tom doesn't do anything close to treating Jerry as a child, while something about the way John deals with his grandfather is infantilizing.

I think the answer lies, not just in the behavior of Tom or John, but in their mental stances towards their interlocutors. I suggest that the missing component of epistemic paternalism is *a lack of respect for the other's epistemic capabilities*.[15] Tom's intervention in no way indicated a lack of respect for Jerry's rational abilities, while John's manipulation of his grandfather did. Thus, an intervention becomes an instance of epistemic paternalism when the intervention is performed without respect for the other's capabilities as a rational epistemic agent.

We need to add a fourth condition to Jackson's account, requiring epistemic paternalism to be performed with a lack of respect for the other's epistemic capabilities. I'll say more about what exactly it means to show (or fail to show) another respect for their epistemic abilities. For now, on the hopeful assumption that this suggestion strikes the reader as roughly correct, I'll move to finish revising Jackson's account. The hope is that this change solves both of our lingering problems in one fell swoop: lack of respect carves cases of epistemic paternalism apart from simply forceful everyday collaboration, and this lack of respect generates the pro tanto moral reason against engaging in epistemic paternalism.[16]

[15] Or 'abilities'—I'll use these terms interchangeably.

[16] If you fail to respect capabilities that someone has, this is a basic, pro tanto moral wrong.

This revision has two parts: we first need to change condition (ii).[17] As stated, (ii) simply states the intervention occurs without the agent's actual consent. But we share evidence, information, and reasons with others all the time without obtaining consent first, and these cases, even if they meet (i) and (iii), do not feel paternalistic (i.e., Tom and Jerry). Additionally, as I pointed out, it's implausible that we have a wide-ranging claim-right against others sharing information with us without first obtaining our consent.

The essence of epistemic paternalism is not *just* that a paternalistic party interferes without the recipient's consent—it's that this interference is, in principle, something to which the other *cannot* consent. This subtle change provides the asymmetry we're after: the paternalistic party will engage in an intervention for another's own epistemic good, whether or not they *do* or *would* consent to this intervention. The recipient's will plays no role in what the paternalistic party does. Thus, being paternalistic in this way fails to show the other respect; it makes them into a means, rather than an end.[18]

Our final revision, then, is to add lack of respect for the other's rational epistemic capabilities as a fourth condition that is necessary to define epistemic paternalism. These revisions to Jackson's definition create the following account:

> **Epistemic paternalism:** (i) intentionally and significantly interfering with someone's inquiry, (ii) in a manner which bypasses the possibility of their consent or dissent, (iii) for what the paternalistic party takes to be their own epistemic good, (iv) while lacking respect for their relevant epistemic capabilities.[19]

Section 2: Refuting Descriptive Skepticism

With this revised account in hand, we can now examine and refute recent skepticism about epistemic paternalism pushed by Pritchard, Bullock, and Jackson.[20] In her recent article, Jackson advances two forms of skepticism; the first denies the existence of the phenomenon of epistemic paternalism. We might call this

---

[17] Which, curiously, is the only condition Jackson doesn't revise.

[18] This revision makes various forms of epistemic paternalism coalesce around a similar wrong-making feature. One could manipulate another's beliefs in all sorts of ways: by lying, playing on the other's emotions, using framing effects or nudges, or by presenting true information in a strategic manner. Under our new account, these various ways of being epistemically paternalistic share the common wrong-making feature of treating the other as a means rather than as an end. See O'Neill (1985) for a discussion of how deception and coercion use another as a mere means in this way.

[19] A few further clarifications may be needed. First, when I say that the intervention shows as lack of respect for the other's *relevant* epistemic capabilities, I say this to indicate that one can engage in epistemic paternalism with respect to some local failing without believing that the other is an entirely incapable epistemic agent. An intervention might target a local failing without implying the agent's abilities have failed globally. For instance, John's grandfather may have fine mathematical skills, but if those aren't relevant to the topic at hand, John doesn't have to disrespect those abilities when he engages in EP. I'm grateful to Yoorim In for suggesting I say this explicitly.

Second, the term 'inquiry' implies that epistemic paternalism can only be undertaken while the other person is actively investigating some question or issue. On a typical understanding of inquiry, suspending judgement about P counts as active inquiry, while forming a belief or securing knowledge that P is not inquiring. So, for instance, double-checking your calculations is active inquiring, and when you finish this process, you've stopped inquiring. Because inquiry is a very active term, I may be pointing to something larger in scope than just inquiry, at least on a standard understanding of that term. Perhaps my account should refer to the interference of someone's inquiry, evidence gathering, and belief maintenance, for instance. My point is that I want our account to capture both interfering with someone's active inquiry, but also cases wherein someone has settled on P and I forcefully jumpstart (and steer) their subsequent inquiry process. But because inquiry is interpreted in such active terms, my fear is that this will misleadingly restrict the scope of the phenomenon. Epistemic paternalism should include interfering with both active and passive doxastic states.

[20] None of my pushback on Jackson's skepticism inherently relies on my revisions to her account of epistemic paternalism.

'descriptive skepticism'—the denial that we do in fact engage in epistemic paternalism. The other form (which we might call 'normative skepticism') denies that engaging in epistemic paternalism is ever all-things-considered appropriate or justified.[21] My target in this section is descriptive skepticism. With precision and clarity, Jackson summarizes her argument as follows:

> When defining epistemic paternalism, we are faced with a dilemma. If epistemic paternalism requires acting only for another's epistemic good, then we're spending a lot of intellectual energy and ink analyzing a practice that rarely occurs (and most of our examples don't actually apply). However, if an action can count epistemic paternalism if it is motivated by both epistemic and non-epistemic reasons, then it becomes less clear what distinguishes epistemic paternalism from regular paternalism, why it would deserve its own literature and analysis, and why its justification would differ from that of regular paternalism (2021, 140-141).[22]

Both prongs of this dilemma share an assumption that the type of paternalism is determined by the motiving reason(s) for that paternalism (2021, 137-139). That is, if we look at an agent's motivation for engaging in paternalism, the domain of that motivating reason determines the type of paternalism they are engaged in. So if a paternalistic agent acts for *moral* considerations, then their paternalism is *moral*. If they have *epistemic* reasons for their intervention, then their paternalism is *epistemic*. And not only must the paternalistic party have *epistemic* motivations for their activity, but their reasons must be *solely* epistemic if the subsequent paternalism is to count as epistemic. Thus, for Jackson, an instance of paternalism counts as epistemic if and only if:

1. the agent has one or more motivating reasons which are epistemic in nature and
2. the agent lacks any motiving reasons which are non-epistemic.

Should we accept condition 1—that the type of paternalism is determined by the motiving reason(s) for that paternalism? We can call this position, pushed by Pritchard, Bullock, and Jackson, the *Determination by Motivation* view: that the reason for one's engaging in an activity determines the kind of activity it is. This idea sounds plausible enough. But when considered further, this way of determining an activity's domain is quite odd. If one's motivating reasons for engaging in ϕ determine what kind of activity ϕ is, we quickly generate conceptually bizarre conclusions. For instance, there was a morally-loaded motivation for the Manhattan project's scientific investigation of nuclear fission (defeating the Nazis and so forth). Does this mean the inquiry into how uranium isotopes work wasn't *scientific*?—i.e., *epistemic*? Likewise, if I have a moral motivation for playing a sport (to better my health), is that sport no longer an *athletic* activity? The *Determination by Motivation* scheme clearly misfires in these cases; surely we can have a moral reason for engaging in an athletic activity.

Similarly, assumption 2 cannot stand. If, for ϕ to be an epistemic activity, I cannot have *any* non-epistemic reasons for ϕ-ing, then the problems discussed in the prior paragraph are only multiplied. On this assumption, if I have multiple motivating reasons for playing a sport (some athletic, some moral, some epistemic: learning a new skill), then my playing this sport is no longer an athletic activity. For my taking a pottery class to be an aesthetic activity, I cannot have any non-aesthetic motivations for taking the class.[23]

---

[21] The two forms of skepticism are, in theory, wholly independent.

[22] In footnote 7 immediately following this quotation, Jackson adds: "One possibility is that both general and epistemic paternalism essentially concern the same phenomenon, but are simply evaluated from different points of view; the former evaluates it morally and practically, and the latter evaluates it epistemically…While I'll proceed as if epistemic paternalism is a unique phenomenon, this suggestion warrants further exploration" (2021). This is not precisely the line of rebuttal I take, but it is similar.

[23] Furthermore, we could apply this same standard to moral (or *non-epistemic*) paternalism. How many instances of interpersonal paternalism wholly exclude any motivation to better the other's epistemic situation? In other words, how much paternalism aims only at manipulating action without any change in doxastic states? Of course, there are some cases, like a parent coercing a young child to eat their vegetables without aiming to create the belief they are healthy.

The straightforward solution is to instead recognize that the type of activity ϕ *is* depends on the nature of ϕ, not one's motivations for engaging in ϕ-ing. We can call this scheme *Determination by Domain*. Applied to our current question, this means the type of paternalism is determined by the domain of the *target* of one's manipulation: if your aim is to manipulate someone's action, then the paternalism is moral or pragmatic (or at least *non*-epistemic), and if your aim is to manipulate someone's doxastic attitude, reasoning, or inquiry process, then the paternalism is epistemic. This makes it the case that I can have a (partially) moral story for why I'm being paternalistic while engaging in epistemic paternalism by manipulating another's beliefs. What identifies epistemic paternalism as such is the target of manipulation being *epistemic* in nature—another's belief or doxastic attitude, reasoning, inquiry process, etc.[24]

Indeed, I suggest that most interpersonal paternalism (at the very least, a great deal of paternalism) is epistemic—that many cases of ordinary paternalism have an epistemic target, aiming to modify another's doxastic states, reasoning, inquiry process, etc. Much of our paternalism aims to manage others' beliefs and subsequent voluntary behavior through manipulation, deception, gaslighting, nudging, etc.[25] These interventions target their subject's epistemic states or processes in an attempt to alter downstream behavior. Especially in modernized cosmopolitan societies, a great deal of paternalism takes this form rather than straightforward coercion, threats, or outright physical control. For instance, if John wants to get his grandfather to vote 'yes' on a proposal to increase funding for climate change, the most natural way to pursue this aim is to manipulate his grandfather's beliefs about the importance of mitigating climate change. This change in belief will then likely instill a change in behavior—at least this is what John hopes for.[26]

Seatbelt laws may be like this as well. These laws don't care about what citizens *believe*; they just care about what they *do*. But these cases are few and far between in interpersonal contexts. Most of the time, when we act paternalistically towards one another, at least part of what we're doing is attempting to change the other's doxastic states.

[24] Here is one possible reply to my argument for *Determination by Domain*: Sports can still be an athletic activity with a mixture of motivating reasons behind engaging in that activity. But there's something special about epistemic phenomena: the epistemic domain itself excludes non-epistemic considerations from ever playing a role. (This is, for example, a widely held evidentialist source of skepticism with respect to pragmatic or moral encroachment on belief and knowledge. The central idea is that practical or moral concerns simply cannot generate *epistemic* reasons for or against belief; these sorts of concerns are the wrong kind of reasons for any changes in epistemic states.) The epistemic domain is special in its need for purity: no non-epistemic concerns get to influence what goes on in the epistemic domain. So if we apply this style of reasoning to epistemic paternalism, perhaps we should also think that it cannot be undertaken for any non-epistemic motivations. I doubt that I can give a decisive reply to this style of reply. I'll only say that this seems rather non-parsimonious, and that I am less concerned with the label that we attach to a phenomenon than I am about the nature and justification of the phenomenon itself. My point is that, contrary to Jackson, we *do* seem to engage in the manipulation and management of each other's doxastic states. (And, as I'll soon argue, this can sometimes be justified.) Whether or not we want to resist the term 'epistemic paternalism' is of less concern to me, though I think it's useful and straightforward to refer to the management of others doxastic states as epistemic paternalism.

[25] On the widely accepted view of doxastic involuntarism, we would expect epistemic paternalism to take the form of nudging, manipulating, lying or deceiving, or gaslighting rather than outright threats or coercion. We cannot coherently threaten others to 'believe P or else' in the same way we can threaten them to 'do ϕ or else.' This is reason to believe that most cases of paternalism which don't involve direct threats are instances of partially epistemic paternalism: attempts to manipulate beliefs or reasoning which then influence what the other voluntarily does on that basis.

[26] I suspect that most everyday person to person paternalism is mixed because we often lack the ability or wherewithal to issue outright threats, coerce, or physically force others to do what we know to be in their own interest. It's often more effective (and less risky) to instead attempt to influence what people do by influencing what they believe.

Even at the level of governmental policy making, where we have the power to force others to act rationally for their own collective wellbeing, it still makes sense to try and paternally encourage others to arrive at these conclusions themselves. In discussing the need to resolve discordance between scientific and public opinion on polarized political topics, Anderson suggests the following: "Another measure to overcome partisan segregation would be to recruit spokespersons of diverse backgrounds to inform the public of the scientific consensus. People tend to accord higher credibility to the testimony of people who share their background and value orientation. President Obama's cultivation of ties to leading Evangelical Christians, such as the influential preacher Rick Warren, who promote action against global

Jackson's descriptive skepticism was based on an unduly restrictive requirement. On her view, epistemic paternalism must be entirely distinct from 'general' paternalism; it must be undertaken solely for epistemic motivations. By setting the bar unreasonably high, skepticism came easily. But with a clearer account in hand, we can see that epistemic paternalism is a widespread phenomenon. It's quite rare for us to threaten or physically force others to undertake certain actions. Instead, I suspect we more often engage in epistemic paternalism, attempting to manage another's beliefs to then affect their downstream behavior. Thus, rather than being exceedingly rare, epistemic paternalism is exceedingly common. Investigating normative questions about epistemic paternalism is, contrary to Jackson's suggestion, well worth our time and ink.

Section 3: Respect for Another's Epistemic Capabilities

I promised to say more about respect for another's epistemic capabilities. What does this amount to? How do we show others respect for their epistemic capabilities? And what does it mean to withdraw this respect?

Most broadly, we respect someone when we acknowledge their agential capabilities and accordingly respond to how they exercise those capabilities.[27] For instance, if a colleague crafts a particularly clever philosophical argument, we respect them by responding accordingly; we recognize the quality of their epistemic performance and feel something like admiration or esteem. We might even verbalize this feeling by saying that we're impressed by the argument or praise them in other ways. In respecting another, we do not treat them in accordance with an instrumental goal, which would render them a mere means to achieve that aim. We do not, for instance, express admiration or praise only as a way to ingratiate ourselves or to curry favor. Doing so would be to treat the other as a mere means—i.e., to use them instrumentally for an ulterior aim—and thus fail to respect them.[28]

We respect someone when we respond to them in accordance with how they exercise their agential capabilities. The details of *how* we show this respect vary depending on the kind of agential capabilities we show respect for. For instance, to respect an athletic opponent is to take their abilities seriously: to treat them like they can beat you, that they are worthy of your efforts to win, etc. This is why we tend to think that letting someone win demonstrates a lack of respect. Of course, there are times when letting someone win might be the right thing to do—but think about how it feels to learn that someone didn't even try in a game of racquetball. It feels disrespectful to be *given* a win when what you wanted was to *compete* for it. By disrespecting your athletic abilities, your opponent won't give you the opportunity to exercise your agency in the right way (by genuinely *winning*.)

---

warming, implements this strategy. Warren is more likely to be able to open the minds of conservative Christians to the realities of global warming than Al Gore. On a larger and more speculative scale, the epistemic powers of democracy might be enhanced by nonpartisan integrative redistricting. Representative districts today are dominated by partisan gerrymandering, which tends to create politically homogenous districts that, in turn, lead to highly polarized politics. If candidates couldn't win elections by only appealing to their own party, they would be less likely to run on dogmatic claims that have little traction outside their party. Evidence taken seriously outside the partisan ingroup would be more salient in political campaigns and less likely to be dismissed" (2011, 158-159). I wholeheartedly agree with these suggestions. But notice the overtly paternalistic stance being taken here: we need to manage certain segments of the US population towards believing the right things, and the strategies being considered here are weighed by how *effective* they might be. We need to get spokespeople who fit the same epistemic demographics to bring those lagging behind up to speed on issues like climate change (which at our current political moment feels like a luxury to worry about).

[27] I'm not taking a stand on whether this is the fundamental form of respect, whether there are other ways of respecting people, etc. I simply mean to say that, in general, treating people as a means—i.e., using them instrumentally—fails to respect them as autonomous agents.

[28] See O'Neill (1985).

In the moral domain, we show others respect for their agential capabilities by holding this same expectation that they can and will exercise their abilities. We hold people to moral expectations by becoming angry, resentful, and indignant with them when they act wrongly (Strawson 1962). This might sound strange at first pass; how could getting angry with someone be a way of showing them respect? While it's certainly possible to express one's anger in a disrespectful way, anger or resentment itself is actually a fundamentally respectful response to another's failings. When we react with resentment or indignation in response to another's moral failures (i.e., a failure to exercise their moral capabilities), we reveal an expectation: we thought they could have done better than they did, and we expected more of them (Sher 2006, Korsgaard 1992 & 1996).[29]

It is the withdrawal of this 'basic demand'—that others exercise their abilities well—which both constitutes and signals a lack of respect for the other's capabilities as a moral agent (Watson 1987). When we lose respect for someone as an agent, we fail to demand or expect that they exercise their normative abilities. Of course, sometimes this withdrawal of respect is brought about for good reasons. As Watson has pointed out, we exempt those with diminished capacities from our interpersonal accountability practices, and we judge this to be the morally right thing to do (1987). So, for instance, when a relative with a brain tumor or severe dementia makes an offensive remark at Thanksgiving dinner, we don't react with anger or blame like we would if our partner had made that same comment. But for those who meet the necessary conditions to be held accountable, being treated as though one was exempted is fundamentally disrespectful. This is why we have phrases like being 'beneath contempt.' Even negative emotions like contempt, anger, resentment, and indignation reveal a kind of respect for their recipient's agential capabilities. These reactions contain an implicit message; they convey the sentiment that their recipient could have and should have exercised their capacities better than they did (Darwall 2006 & 2024).

We can respect (or fail to respect) another's epistemic capabilities in this same manner. When we respond according to the quality of another's epistemic performance, we show them respect for their capabilities. When we genuinely praise someone for their incisive argument, for instance, our admiration is a response to the quality of their performance. Conversely, being led like a lamb to a conclusion, as John does to his grandfather, is disrespectful in the same way; John does not respect his grandfather's agential capabilities in the epistemic domain. John does not treat his grandfather as someone who can handle his full thoughts on the matter. He must be managed towards the right conclusion from within a broken epistemic framework, one in which the only way to justify caring about the Earth is to relate this concern to some matter of religious importance. John's mental stance involves a withdrawal of a basic expectation that his grandfather be able to work through the matter, handle a full account of why one should invest in climate change, etc. When John engages in epistemic paternalism towards his grandfather, he withdraws from open, reciprocal engagement to an instrumental view of his grandfather, managing his inquiry rather than naturally communicating with him.

Importantly, showing respect for another's epistemic abilities doesn't always have to be emotionally toned; sometimes we show someone respect simply by honestly (and maybe even forcefully) disagreeing with them. Sometimes, showing respect takes the form of saying "I disagree with you about that." This is precisely the way in which John withholds respect from his grandfather; he distinctly does not disagree with him in an open, honest manner. Instead, he works from within what he considers to be a broken epistemic framework and attempts to manipulate or manage his grandfather towards the right conclusion.

Thus, to respect another's epistemic capabilities is to react in a manner which accords with the quality of their exercise of these capabilities. To show this respect is to respond on this basis, expressing one's inner

[29] Strawson (1962, 23-24; 29-30) refers to this expectation as what has become known by the 'basic demand' for goodwill (or at least the lack of ill will). I believe the term 'basic demand' originates from Watson (1987).

reaction simply as a response to how the other performed; to treat the other as an end rather than a mere means.[30]

I'll now take a moment to pause and offer a summary of what we've done thus far. First, we developed a revised account of epistemic paternalism; it's an active way of engaging with another constituted by (i) intentionally and significantly interfering with someone's inquiry, (ii) in a manner which bypasses the possibility of their consent or dissent, (iii) for what the paternalistic party takes to be their own epistemic good, (iv) while lacking respect for their epistemic capabilities. We then addressed Jackson's descriptive skepticism about the existence of epistemic paternalism. In this section, I said more about the type of respect that's withdrawn when engaging in EP. At this point, we have an account of what epistemic paternalism is, a defense of its widespread existence, and a deeper understanding of the constitutive loss of respect for the paternalized party. In the next section, I'll move to my second aim of the paper: using this revised account to architect a framework for answering the pressing normative questions regarding epistemic paternalism.

Section 4: Normative Questions

This section has three aims. First, I want to give a partial explanation for why Section 1 mattered; I want to explain why our account of epistemic paternalism affects how we go about answering normative questions. This will involve showing why Jackson's account unfairly stacks the deck in favor of normative skepticism about the appropriateness of epistemic paternalism. Second, I want to provide a more positive framework for answering normative questions. Third, I want to give a high-level answer to one such question.

Let's begin by examining how our descriptive account can affect our normative analysis (i.e., whether epistemic paternalism is or can be justified). Here I'll return to Jackson's framework; she first provides four questions about (epistemic) paternalism's normative status:

Q1. Is paternalism all-things-considered justified?

Q2. Is paternalism epistemically justified?

Q3. Is epistemic paternalism all-things-considered justified?

Q4. Is epistemic paternalism epistemically justified?[31]

First, this already strikes me as a somewhat flat-footed way to frame an investigation of whether (epistemic) paternalism could be justified. When we ask '*Is* epistemic paternalism justified?,' this frames the question in a way that implies the answer is necessarily all-or-nothing. Either it *is* or *is not* justified. But there are more fine-grained distinctions available. For instance, if we agree with Jackson that paternalism disrespects another's autonomy, then this is a moral reason against engaging in paternalism (of any kind). But plausibly, in some cases, other goods can justify this pro tanto moral wrong. So I'm skeptical of this way of framing Q1-Q4.[32] It would make more sense to ask (i) What exactly is the moral wrong in (epistemic) paternalism? Then

[30] This does not mean that every mistake in the epistemic realm has to be met with opprobrium or condemnation, just like not every mistake in the moral domain has to be met with anger and resentment. For instance, when a colleague fails to consider a certain snag in their philosophical argument, it would be unfitting to respond this way. But this doesn't mean that you fail to respect their epistemic abilities; it simply means that crafting an airtight philosophical argument is *extremely hard*, and mistakes are par for the course. So there's nothing to fittingly feel opprobrium towards.

[31] Jackson (2021, 142).

[32] This is a relatively small point and admittedly whiny point. But how we frame the questions at hand matters because their framing influences the range of answers we can coherently provide. The question 'When did you see the woman at the crosswalk?' already constrains the range of congruent answers. Replying that you did not see a woman at the crosswalk in an implicit way of reframing the question at hand. We need to do the same here.

ask (ii) Under what conditions could (epistemic) paternalism could be all-things-considered justified? The first question was answered in Sections 1 and 3; the latter question will be explored here.

I worry that Jackson's account of what epistemic paternalism *is* stacks the deck against whether it could ever be *justified*. Recall that Jackson determined the type of paternalism by reference to the type of reasons for its use, regardless of the interference targeted. I have already argued that this is an unpromising descriptive framework, but now I can spell out how this error derivatively pollutes Jackson's later analysis of epistemic paternalism's normative status. Consider the following suggestion:

> Several of those who have written on the justification of epistemic paternalism thus far have focused on Q3. However, this strikes me as a debate that isn't especially fruitful, because it seems difficult, if not impossible, to justify epistemic paternalism on all-things-considered grounds…whatever epistemic gain supposedly motivates epistemic paternalism—e.g. true belief, justified belief, knowledge—will almost always be outweighed by moral considerations in favor of personal autonomy and sovereignty. (Recall here that, if we understand epistemic paternalism as primarily motivated by epistemic considerations, we cannot count the moral or practical value conferred by the content of the beliefs; it is purely the epistemic value of the true/justified beliefs/knowledge) (Jackson 2021, 142).

On the following page, Jackson summarizes this thought as follows: "In general, it is hard to see how epistemic paternalism could be morally or all-things-considered justified, especially if it requires acting *solely* for another's epistemic good" (2021, 143 emphasis added).

Right. But this conclusion is baked into the premises. If epistemic paternalism requires that the paternalistic party has no non-epistemic motivations for engaging in that paternalism, then of course it's going to be nearly impossible to justify epistemic paternalism on an all-things-considered basis, since we also know that engaging in paternalism of any kind disrespects another's agency. So we've set up the question of whether epistemic paternalism could be all-things-considered justified as essentially the question: *could a small, purely epistemic good ever outweigh a serious moral wrong?* Of course it's then easy to answer this question negatively.

But this conclusion only works if one's definition of epistemic paternalism excludes the possibility of managing another's doxastic states or processes for epistemic, practical, or moral considerations. In other words, I'm suggesting that Jackson's conceptual scheme makes it definitionally impossible to justify epistemic paternalism. Say I'm considering whether to manage my grandfather's credence about P. I need to figure out if this paternalism is all-things-considered justified. But if I appeal to any moral reasons for engaging in this epistemic paternalism, Jackson will reply that, because I have both epistemic and moral reasons for my interference, it's no longer *epistemic* paternalism (even if I'm solely manipulating his doxastic state). Jackson has asked us to provide an all-things-considered justification, while covertly making it impossible to provide an *all*-things-considered justification. This is unfair.

The analogous scenario would be me asking you to justify spending any resources on an athletic facility (a field, stadium, etc.) while there are moral needs that could be met instead, while also defining the type of facility by the type of reasons one has for its construction/maintenance. So if you give me any moral reasons to build a soccer field (it will bring the community together, create a shared space, provide an opportunity for kids to exercise and make friends), I'll say it's now a *moral* facility, not an *athletic* one. But this way of analyzing the question is profoundly unhelpful. I would unfairly make it definitionally impossible for you to succeed in all-things-justifying an athletic facility.

We need a new framework for considering normative questions about epistemic paternalism. This is my second aim of the current section.

Recall one of our conceptual revisions to Jackson's account. In her original account, it was not clear where exactly the moral wrong of epistemic paternalism originated. We solved this issue by requiring epistemic paternalism to involve a lack of respect for the other party's epistemic capabilities. (This explained why John's behavior towards his grandfather was paternalistic while Tom's intervention of Jerry's inquiry was not.) This feature provided the pro tanto moral wrong of epistemic paternalism.

But notice that this feature not only clarifies our concept—it also helps adjudicate when epistemic paternalism is justified. By identifying the wrong-making feature—treating another as though they lack capabilities they in fact possess—we can identify that which must be overcome if an instance of epistemic paternalism is to be justified. Thus, we can move to the third and final aim of this section and identify three types of general cases in which epistemic paternalism is all-things-considered justified.[33]

First, engaging in epistemic paternalism is justified when the other person is in fact lacking the relevant epistemic capabilities. This condition offers us a straightforward explanation; if the moral wrong of epistemic paternalism results from disrespecting another's capabilities, then if that person lacks these capabilities, the wrong-making feature is absent, and EP is therefore permissible. We can see this idea in action by looking at the semantic root of paternalism; in essence, engaging in paternalism is treating another like a child by acting like their parent. But when that person *is* a child, there's nothing wrong with being paternalistic. Parents need to restrict what their eight-year-old can do for their own good.[34] Tension arises when we treat other adults in this way precisely *because* adults typically possess the capacities which paternalism disrespects.

But of course, sometimes other adults do lack these capabilities. How do we know when this condition holds? One clue can be found by examining the moral responsibility literature's treatment on exemptions. In brief, exemptions are certain facts about a person's constitution which render them inapt for accountability practices—blaming, punishing on a retributive basis, etc. There are broadly agreed upon conditions which fit this bill: a brain tumor in the right areas, abusive upbringing, mental illness or disease, etc. These conditions make a person an unfitting recipient of reactive attitudes like anger, opprobrium, and resentment precisely *because* their agential capacities have been sufficiently diminished.[35] A plausible line of thought is that whenever a person is exempted from accountability attitudes or practices, they are also permissible targets of epistemic paternalism, since those agential capacities required to be aptly held accountable are also those whose absence remove the wrong-making feature of paternalism.[36]

Second, epistemic paternalism is justified when its benefits outweigh the moral wrong of disrespecting another's autonomy. I take the generality of this second case as a virtue; it leaves open what benefits (moral, epistemic, or otherwise) would in fact outweigh the moral wrong of EP. This formulation leaves the question open in two senses. First, it makes no judgement on whether the relevant benefits must be moral or epistemic (or otherwise).[37] Second, it makes no judgement on what magnitude of benefits are required to outweigh the wrong of disrespect. Thus, a hard-nosed Kantian can reply that no amount of goods can justify the disrespect

[33] Justified here should be seen in the sense that it renders epistemic paternalism permissible rather than obligatory. Additionally, I am not claiming these are the *only* three justified cases.

[34] This relationship explains why parents become less overtly paternalistic as their children grow; it's because they develop greater capabilities.

[35] While excuses aim to change the nature of an action, exemptions aim to cancel blame by rendering the *person herself* an inappropriate subject or recipient of blame (Strawson 1962, Watson 1987, Sliwa 2019, Kauppinen 2026).

[36] And, as O'Neill rightly points out, these cases are made more complicated when we consider the fact that even typically rational adult agent can become exempted only for a short time: "Paternalism may then seem permissible, even required, for those who are, if temporarily, *only* patients…Even when we are mature we are seldom ideal rational patients" (1985, 256).

[37] A move I can make but Jackson cannot.

involved in using someone as a mere means, while a less deontic view can posit cases on both sides. This second set of cases provides a general framework to evaluate this question with greater clarity.

Third, and again only generally, epistemic paternalism is plausibly justified in any case where non-epistemic paternalism is justified. By non-epistemic paternalism, I mean the same sort of intentional and significant interference for a person's own good which does not have as its target the other's epistemic states or processes (beliefs, reasoning processes, inquiry, etc.). This is paternalistic intervention which does not attempt to change behavior by changing what one believes, but merely attempts to forcibly change behavior by coercion, threats, physical control, etc.

For instance, consider a case wherein non-epistemic paternalism seems clearly justified: a Seventh Day Adventist husband and wife refuse to allow a life-saving blood transfusion for their child. The hospital, I submit, would be doing the right thing to engage in a paternalistic override of the parents' decision. This instance of non-epistemic paternalism would be justified. (It's non-epistemic because the interference does not aim to modify any of their doxastic states or inquiry processes; it simply overrides what the parents allow to be *done* to their child.) But if this interference is justified, then it's plausible that epistemic paternalism would also be justified in this case. Succeeding in this aim would be to change the parents' relevant doxastic states (and therefore their subsequent decision). Perhaps the hospital's medical ethicist could accomplish this by persuading the parents of a more liberal interpretation of Genesis, Leviticus, and Acts. Of course, the medical ethicist could believe that determining what health care is permitted by examining the Bible is an absurd epistemic system to live by, but this would not prevent her from manipulating the parents' beliefs for their own good—and more importantly, their child's. In fact, successful epistemic paternalism might be a vastly better result than simply compelling action (or inaction), since the parents' change in belief might spread to other members of their family or religious community.

This is, of course, only a rough rule of thumb, and I have no doubt that edge cases can be found. In some instances, manipulating others' beliefs, even for their own good, might be worse than compelling them to act. But there are clear reasons to accept this general scheme. For instance, if you would be justified in physically restraining someone, it's plausible to think you'd also be justified in deceiving him so that he 'restrains himself' voluntarily. This would be to accomplish the same aim by nonviolent (or at least non-physical) means. Thus, in any case where non-epistemic paternalism is justified, it's plausible to think that epistemic paternalism is justified as well.

Section 5: Epistemic Paternalism and Strawsonian Objectivity

Let us pause to take stock of what we've said thus far: epistemic paternalism is an active way of engaging another with the aim of modifying their doxastic state or inquiry process. Because the intervention will occur regardless of what the other wills, they cannot consent or dissent to the intervention in principle. Engaging in epistemic paternalism thus involves a lack of respect for the other's relevant epistemic capabilities; perhaps they can't be trusted to find and weigh evidence on their own or to put the pieces together themselves. They are not treated as an independently capable epistemic agent; instead, the other is used as a mere means for the end of modifying their own doxastic states or processes.

In the preceding sections, many readers will have heard the whisper of a Strawsonian theme. I suggest that epistemic paternalism is a way of engaging another that relies on taking up a certain kind of mental perspective: what we might call an epistemic objective stance. I have in mind the same type of objective attitude discussed in the Strawsonian strain of moral psychology. In the broadest of strokes, to inhabit an objective

stance towards someone is to see or regard them in a way which renders them closer to a mere *thing* rather than another *agent*.[38] In a famous line, Strawson introduced the notion of objectivity as follows:

> To adopt the objective attitude to another human being is to see him, perhaps, as an object of social policy; as a subject for what, in a wide range of sense, might be called treatment; as something certainly to be taken account, perhaps precautionary account, of; to be managed or handled or cured or trained …[objectivity] cannot include the range of reactive feelings and attitudes which belong to involvement or participation with others in inter-personal human relationships; it cannot include resentment, gratitude, forgiveness, anger, or the sort of love which two adults can sometimes be said to feel reciprocally, for each other. If your attitude towards someone is wholly objective, then though you may fight him, you cannot quarrel with him, and though you may talk to him, even negotiate with him, you cannot reason with him. You can at most pretend to quarrel, or to reason, with him (1962, 25).

The final lines begin to approach the point I have in mind; if I am wholeheartedly engaging in epistemic paternalism, then while I can speak to another in certain ways, I cannot argue or converse with them (though I can still attempt to manage their doxastic states). In a way, this is to accept another famous line of Twain's advice: '*Never argue with stupid people. They will drag you down to their level and then beat you with experience.*' When Twain admonishes us not to *argue* with the less-epistemically-endowed, he's not suggesting we never speak to them. But on one way of interpreting this advice, we're being told to objectively 'work on' certain people's doxastic states rather than engage with them as epistemic equals—in Strawson's terms, to simply manage, handle, or cure some epistemic failings rather than quarrel with someone who lacks the capacity to be *reasoned* with as an equal.

I am suggesting that the *activity* of engaging in epistemic paternalism involves a quintessential *mental state*: a Strawsonian objective stance adopted with respect to the other's epistemic failings. So what does adopting an objective stance involve? Elsewhere, I have argued that Strawson's objective stance is a mental perspective (a way of seeing or regarding someone) that is characterized by (i) an intentional lack of emotional and attitudinal vulnerability, (ii) creating or maintaining a mental distance from the other, and (iii) a proneness towards instrumental thinking about the other.[39]

First, making yourself emotionally or attitudinally vulnerable gives someone the power to affect your inner mental life. Taking up an objective stance limits the other's power to affect your inner life in this way. So while an insult might typically engender anger and resentment, if I've successfully taken up an objective view, I'll instead remain emotionally inert in the face of a derogatory comment. I won't *react* in the manner I typically would. Importantly, objectivity's limitation of attitudinal access is more involved than simply failing to experience reactive attitudes towards another. I don't feel anything at all towards a stranger on the train who hasn't interacted with me in any way. But that doesn't mean I've taken an objective stance towards him. Adopting an objective view is to *intentionally* limit or eliminate one's proneness to feel reactive attitudes in response to the behavior of another.

In the epistemic domain, this objectivity might be as simple as forswearing certain emotional reactions we might have to other's poor epistemic performances. I have often felt opprobrium and contempt, for instance, when hearing inane and repugnant political beliefs.[40] To take an objective stance would be to limit my propensity to react with these emotions. But epistemic objectivity might also extend to other less emotionally-

[38] As Schroeder has more recently argued, "To be treated as a thing is to be minimized, rather than engaged with, predicted and controlled rather than reasoned with, written off as the product of our environment rather than appreciated for our unique contributions" (2019, 95).
[39] Myers (Work in Progress).
[40] See, for instance, Myers (2025).

toned attitudes like doubt, curiosity, or even a willingness to share one's own thoughts on a matter. These attitudes can also be restricted by adopting an objective stance. For instance, when an Evangelical relative tells me that blasphemy is a sin, I'm not particularly bothered by that claim, but I'm also not particularly curious about why they think so. And I'm not prone to respond by voicing my true thoughts either. Instead, their claim appears as something that doesn't need to be engaged with; it's just something that occurred.

Second, the objective stance holds its subject at a mental distance or remove. Just as a disposition to turn *towards* another is characteristic of reactive attitudes, a disposition to move or remain *away* characterizes objective attitudes (Darwall 2024, 31-32). Rather drawing close to another, adopting an objective stance aims to create or sustain distance from its subject. Strawson speaks of objective attitudes having a "detached" view of their subject, invoking this image of mental separation when suggesting that to adopt an objective stance towards someone removes him from one's community of respected equals (1962, 21; 30; 36).

I suspect we mentally distance ourselves from others, not just in the moral domain, but in the epistemic as well. Strawson speaks of mentally exiling another from one's 'moral community,' but it seems we also do the same for our 'epistemic communities.' When I take an objective stance towards my Evangelical relative, I'm exiling them from my epistemic community of equals—those whose claims I feel I must take seriously. But what exactly does it mean to 'keep someone close' or 'push them away' in the epistemic domain? Roughly, I have in mind something like Gibbard's notion of *thinking cooperatively*: "treating each other's thoughts like thoughts that occur to oneself, to be considered and supported or refuted" (2002, 74). When we think cooperatively as co-members of the same epistemic community, "I try to get you to help me with my thinking, to join with me in thinking" (Gibbard 2003, 275). Thinking cooperatively is to keep another close, to draw them in. By adopting an objective stance, however, one switches modes to something like Gibbard's *judgement individualism*: "On this individualist picture, [two people] may still profit from sharing their plans with each other. Each stands as an island of judgment as to what to do, but each can use the other in limited ways: as a source of testimony, as a proposer of thought experiments, and as a sourcebook of arguments to contemplate and accept or reject. Neither, though, is to treat the other's judgments as his own" (2003, 272). The image of an isolated island illustrates the kind of mental distance involved in taking an objective stance towards others, both as moral and epistemic agents.[41]

Finally, adopting an objective stance makes one prone to view the other in instrumental terms—as a means to some end. Strawson gives us several examples of such ends: we might adopt an objective view "as a refuge, say, from the strains of involvement; or as an aid to policy; or simply out of intellectual curiosity" (1962, 25). We can take up an objective stance to create mental distance from another which then allows for the pursuit of some ulterior aim involving them.

In the epistemic domain, these ulterior aims often involve others' doxastic states—the management of their beliefs, credence, or interrogative attitudes towards certain propositions. As we've seen, this instrumental management involves communicating what is most effective rather than one's natural responses. When objectively managing another, there's a strategic filtering of outgoing communication guided by one's goal—to change someone's belief, to inject doubt or lower the other's confidence, to engender greater openness to a

[41] In response to sustained disagreement, Gibbard suggests two objectively toned ways to move forward: "Suppose, though, I just find you unhelpful and hindering. Then I can respond in at least two different ways. I can simply disagree with much of what you say, rejecting it and trying not to let it influence my thoughts. Alternatively, I may not treat you as even voicing thoughts that I can accept or reject. You have your plans and I have mine, and difference in plans is no kind of disagreement" (Gibbard 2003, 280).

proposition, etc. Recall the way John worked from within his grandfather's broken epistemic framework; he said only what he thought would be *useful*, not what he actually thought.[42]

Interestingly, as Strawson has suggested, there are various ways to both enter and inhabit an objective view. We might actively adopt an objective stance towards a difficult person, while we often find ourselves passively sliding into an objective view of a young child's poor behavior. In our ordinary adult relationships, we most often utilize objectivity as a practical, psychological tool. We often choose to utilize this resource "as a refuge, say, from the strains of involvement; or as an aid to policy; or simply out of intellectual curiosity" (Strawson 1962, 25).

It is this active adoption and instrumental management of another which, I suggest, characterizes the kind of objective mental perspective involved in epistemic paternalism. I argue that engaging in EP involves mentally shutting down one's natural reactions, creating distance to observe and understand the other more clearly, and then using this vantage point to engage in instrumental, strategic use of the other to modify their own doxastic states or processes. Engaging in epistemic paternalism highlights a quintessential mode of epistemic objectivity. Thus, we have the following conceptual scheme: Strawson's objective stance is a mental perspective we can take towards other people. We can passively fall into this stance, or we can actively adopt it for our own practical aims. I suggest this objective view is something we can and do employ, not only in the moral domain, but in the epistemic as well. And epistemic paternalism's *active* aim of managing or manipulating others' doxastic states makes it an instance of an intentional, active objective perspective—just the sort of thing Strawson alluded to when suggesting that, under an objective stance "you may talk to [another], even negotiate with him, [but] you cannot reason with him" (1962, 25).

We can see how this notion of objectivity functions in the epistemic domain by returning to my introductory case. My future mother-in-law's obsession with conspiracy theories has become more than an inconvenience. Until this point, I've cultivated an intentional but fairly passive objective attitude towards her. I simply don't react or engage with her rants and try to steer the conversation away from any lively topics when they arise. I adopt an objective stance as a 'refuge' from the 'strains of involvement,' to use two of Strawson's terms.

But I would need to shift my perspective to a more active, strategic type of objectivity if I were to attempt to manage or repair her epistemic framework—if I were to engage in epistemic paternalism towards her. One tactic I might pursue would be to flood her with even more elaborate conspiracy theories in an effort to push her towards declaring epistemic bankruptcy (as something like a first step towards recognizing that she is in an epistemic community that has a misfiring disposition to read everything as a conspiracy). When she voices skepticism about humans having ever been to the moon, I could outflank her, expressing shock at her naiveté to even believe the moon itself is real. Of course, this would be a *tactic*—a way of pushing her towards seeing that *everything* can be shown in a conspiratorial light.

But in pursuing this strategy, I wouldn't be reporting my actual thoughts on the matter. I don't think it's instructive or worthwhile to spin up extravagant theories about the moon being a hologram. I wouldn't be reporting my unfiltered thoughts on the matter. As Strawson says, while I would be speaking to her, I wouldn't be able to reason or argue with her in a natural, reciprocal manner. I'd be treating her as a means, rather than an end in herself, using her own obsession with conspiratorial thinking against her in an attempt to make her 'bottom out' and declare epistemic bankruptcy.

[42] Consider lying: When we lie to someone, we engage in this same instrumental thinking about what to say in order to produce the desired belief in them. I'm suggesting this same type of thinking applies to paternally managing others' doxastic states in ways that needn't involve outright lies.

In short, when engaging in epistemic paternalism, I switch from saying what actually makes the most sense (to me) to instead saying what I think will be the most *effective*. For instance, when my mother-in-law cites a recent Candace Owens podcast as a 'reputable source,' I might point out all the incentive structures pushing Owens to be bombastic and entertaining rather than scrupulous and careful. But of course, this is an ad hominem attack, and it's not how I'd actually think about the topic myself. But I think that for her, undercutting the credibility of someone like Owens would be a more effective route to dismantling the echo chamber (rather than preaching the values of parsimony or doing a deep dive into any specific claim). Or perhaps I might paint myself as being more open-minded to these ideas than I really am. I might try to paint myself as conspiracy-curious as a tactic to make her believe I was at least on the edges of her epistemic community. Once I had gained that status (and avoided being branded as an outsider), I might then begin to voice doubts about these theories.[43]

These instrumental ways of managing others' doxastic states highlight the kind of active objective stance which I've suggested is the characteristic mental state of epistemic paternalism. To clarify, when I suggest that there is an epistemic objective stance, I am not arguing that there is, in addition to a moral objective stance, *also* an epistemic one that is separate phenomenon. I'm arguing that the objective stance has extended beyond the moral domain from the start.[44] Our ability to adopt an objective attitude towards someone for our own practical aims has always been something we do in response to both poor moral and epistemic agents. In other words, we can take up an objective stance as a tool of management and refuge in response to others' epistemic failings in the same way we can in response to their moral failings.

Conclusion

In this paper, I've attempted to shed light on the descriptive and normative status of epistemic paternalism. First, I developed a refined account of epistemic paternalism aimed at solving two problems. As things stood, a popular conception of EP struggled to differentiate simply forceful collaboration from overtly paternalistic management of another's doxastic states. This previous account also lacked a wrong-making feature to provide the pro tanto moral wrong of epistemic paternalism. These issues were remedied by requiring EP to involve a constitutive lack of respect for the other's epistemic capabilities. With this conceptual clarity in hand, I then defended the descriptive existence of EP by showing that Jackson's *Determination by Motivation* view was fatally flawed and instead defended *Determination by Domain*. Thus, the type of paternalism is determined by the domain of the target of that paternalism, rather than one's motivations. I then developed a rudimentary framework for answering normative questions and suggested three broad classes of cases in which epistemic paternalism might be all-things-justified. Finally, I proposed that the paradigmatic mental state involved in epistemic paternalism illustrates an active form of Strawsonian objectivity.

[43] Similarly, if you know that someone loves to be a contrarian, you might strategically say something close to the opposite of what you actually think in order to get them to believe the truth.

[44] To be precise, we might speak of 'an objective stance in the epistemic domain' rather than a fundamentally '*epistemic* objective stance.' Because the former is semantically unwieldly, however, I use the latter.

References

Alexander, Larry (2014). The Ontology of Consent. *Analytic Philosophy* 55 (1):102-113.

Ahlstrom-Vij, K. (2013). *Epistemic paternalism: A defence*. London: Palgrave Macmillan.

Anderson, Elizabeth (1999). What is the point of equality. *Ethics* 109 (2):287-337.

Anderson, Elizabeth (2011). Democracy, Public Policy, and Lay Assessments of Scientific Testimony. *Episteme* 8 (2):144-164.

Bansback, N. L., Li, L. C., Lynd, L., & Bryan, S. (2014). Exploiting order effects to improve the quality of decisions. *Patient Education and Counseling*, 96(2), 197–203.

Bovens, L. (2008). "The ethics of nudge." In Hansson, M. J. & Grüne-Yanoff, T. (eds.), Preference Change: Approaches from Philosophy, Economics and Psychology., pp. 207–20. Berlin: Springer, Theory and Decision Library A.

Bullock, E. C. (2016). Mandatory disclosure and medical paternalism. *Ethical Theory and Moral Practice*, 19(2), 409–424.

Bullock, E. C. (2018). Knowing and not knowing for your own good: The limits of epistemic paternalism. *Journal of Applied Philosophy*, 35(2), 433–447.

Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press. https://doi.org/10.2307/j.ctv1bzfp0f

Darwall, S. (2024). *The Heart and its Attitudes*. Oxford: Oxford University Press.

Dougherty, Tom (2021). Sexual Misconduct on a Scale: Gravity, Coercion, and Consent. *Ethics* 131 (2):319-344.

Dworkin, Gerald, "Paternalism", *The Stanford Encyclopedia of Philosophy* (Fall 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2020/entries/paternalism/>.

Gibbard, Allan. (2002). The Reasons of a Living Being. In Shafer-Landau and Cuneo (eds.) (2006). *Foundations of Ethics: An Anthology*, Malden, MA: Wiley-Blackwell, pp. 71–78.

Gibbard, Allan (2003). *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.

Goldman, A. I. (1991). Epistemic paternalism: Communication control in law and society. *The Journal of Philosophy*, 88(3), 113–131.

Grill, K., & Hanna, J. (2018). The Routledge handbook of the philosophy of paternalism. New York: Routledge.

Jackson, E. (2020). Epistemic paternalism, epistemic permissivism, and stand point epistemology. In A. Bernal & G. Axtell (Eds.), *Epistemic paternalism reconsidered: Conceptions, justifications, and implications* (pp. 201–215). Lanham, MD: Rowman & Littlefield.

Jackson, Elizabeth (2021). What's Epistemic About Epistemic Paternalism? In Jonathan Matheson & Kirk Lougheed, *Epistemic Autonomy*. New York, NY: Routledge. pp. 132–150.

Jamieson, K. H. and Cappella, J. N. (2008). *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. Oxford: Oxford University Press.

Kauppinen, A. (2026). How to defend yourself against epistemic blame. *Synthese* **207**, 26. https://doi.org/10.1007/s11229-025-05384-x

Korsgaard, C. M. (1992). Creating the Kingdom of Ends: Reciprocity and Responsibility in Personal Relations. *Philosophical Perspectives*, 6, 305–332. Retrieved from http://jstor.org/stable/2214250

Korsgaard, C. M. (1996). *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.

Larmore, Charles (1990). Political Liberalism. *Political Theory* 18 (3):339-360.

Levy, Neil. 2019. "Is Conspiracy Theorising Irrational?" *Social Epistemology Review and Reply Collective* 10 (8): 65-76. https://wp.me/p1Bfg0-4wW.

Levy, Neil (2021). Nudging Well. In *Bad Beliefs: Why They Happen to Good People*. Oxford, GB: Oxford University Press. pp. 132-148.

Liberto, Hallie (2021). Coercion, Consent, and the Mechanistic Question. *Ethics* 131 (2):210-245.

Lougheed, K. (2021). Epistemic paternalism, open group inquiry and religious knowledge. Res Philosophica, 98(2), 261–281.

Manson, Neil (2016). Permissive consent: a robust reason-changing account. *Philosophical Studies* 173 (12):3317-3334.

Mill, J. S. (1869). On liberty. London: Longman, Roberts & Green.

Myers, Jordan (2025). Resultant luck and the target of epistemic blame. *Synthese* 205 (6):1-20.

Nagel, Thomas (1987). Moral conflict and political legitimacy. *Philosophy and Public Affairs* 16 (3):215-240.

Nguyen, C. Thi (2020). Echo chambers and epistemic bubbles. *Episteme* 17 (2):141-161.

O'Neill, Onora (1985). Between consenting adults. *Philosophy and Public Affairs* 14 (3):252-277.

Owens, David (2011). The possibility of consent. *Ratio* 24 (4):402-421.

Pritchard, D. (2013). Epistemic paternalism and epistemic value. *Philosophical Inquiries*, 1(2), 9–37.

Saghai, Y. (2013). "Salvaging the concept of nudge," Journal of Medical Ethics, 39/8: 487–93.

Schroeder, Mark (2019). Persons as Things. In Mark Timmons, *Oxford Studies in Normative Ethics Volume 9*. Oxford, GB: Oxford University Press. pp. 95-115.

Sher, George (2006). *In Praise of Blame*. New York, US: OUP USA.

Simpson, Robert Mark (2013). Epistemic Peerhood and the Epistemology of Disagreement. *Philosophical Studies* 164 (2):561-577.

Sliwa, P. (2019), The Power of Excuses. *Philosophy & Public Affairs,* 47: 37-71. https://doi.org/10.1111/papa.12139

Strawson, P.F. (1962). "Freedom and Resentment." In Russell, Paul & McKenna, Michael (eds.) (2006). *Free Will and Reactive Attitudes: Perspectives on P. F. Strawson's 'Freedom and Resentment'*. New York, NY, USA: Routledge.

Tadros, Victor (2016). *Wrongs and crimes*. Oxford, United Kingdom: Oxford University Press.

Watson, Gary (1987). 4. Responsibility and the Limits of Evil: Variations on a Strawsonian Theme. In John Martin Fischer & Mark Ravizza, *Perspectives on moral responsibility*. (1993) Ithaca, NY: Cornell University Press. pp. 119-148.

Wilkinson, T. M. (2013). "Thinking harder about Nudges," Journal of Medical Ethics, 39/8: 486–486.